

ARTIFICIAL SEMANTICALLY CLOSED OBJECTS

LUIS MATEUS ROCHA

Los Alamos National Laboratory, MS P990, Los Alamos, NM 87545, USA

e-mail: rocha@lanl.gov

Communication and Cognition - Artificial Intelligence. Vol. 12, nos. 1-2, pp. 63-90. Special Issue: *Self-Reference in Biological and Cognitive Systems*. L. Rocha (Ed.)

ABSTRACT

The notion of computability and the Church -Turing Thesis are discussed in order to establish what is and what is not a computational process. Pattee's *Semantic Closure Principle* is taken as the driving idea for building non-computational models of complex systems that avoid the reductionist descent into meaningless component analysis. A slight expansion of Von Foerster's *Cognitive Tiles* are then presented as elements observing Semantic Closure and used to model processes at the level of the cell; as a result, a model of a rate-dependent and memory empowered neuron is proposed in the construction of more complex Artificial Neural Networks, where neurons are temporal pattern recognition processors, rather than timeless and memoryless boolean switches.

1 - Computability and the Church-Turing Thesis.

Within Mathematics computation or computability are notions unequivocally defined by an array of equivalent descriptions such as Turing machines [Turing, 1936], lambda-calculus (equivalent to general recursive functions [Church, 1936]), the Post [1965] symbol manipulating system, Markov [1951] algorithms, Minsky's [1967] register machines, etc.

The objective of formally describing computability lies in defining those functions¹ for

¹ Let S be the set of all objects considered, or universal set. A **Function** with n arguments (or: n -ary function, or: n -ary mapping, or n -ary operation) is a subset of S^{n+1} : $S_1 \times S_2 \times \dots \times S_{n+1}$, with $S_{i=1,2,\dots,n+1} \subseteq S$, $\forall_{i,n \in \mathbb{N}} i \leq n+1$, such that for any two different of its elements $a=(a_1, a_2, \dots, a_{n+1})$ and $b=(b_1, b_2, \dots, b_{n+1})$, we have not simultaneously $a_1=b_1, a_2=b_2, \dots, a_n=b_n$ and $a_{n+1} \neq b_{n+1}$.

The **domain** of a n -ary function f is the set $D_f \subseteq S^n$:

$$D_f = \{(a_1, a_2, \dots, a_n) \mid \exists a_{n+1} \text{ such that } (a_1, a_2, \dots, a_n, a_{n+1}) \in f\},$$

and the **range** (or *codomain*, or *image*) is the set $R_f \subseteq S$:

$$R_f = \{a_{n+1} \mid \exists a_1, a_2, \dots, a_n \text{ such that } (a_1, a_2, \dots, a_n, a_{n+1}) \in f\};$$

Let this relationship be expressed by $f: D_f \rightarrow R_f$. f maps D_f onto R_f , and maps D_f into R , where $R_f \subseteq R \subseteq S$.

Further notice that a function is a special case of a binary relation where each element of the domain can only be related to one element in the range, in other words, one-to-many relations are not functions.

which a **rule** (or *constructive procedure*, or *algorithm*, or *program*) allowing us to obtain range from domain exists - the set of **computable functions**. A function is considered computable if there exists a Turing Machine (or equivalent) defining it, and uncomputable otherwise.

However, proving the computability of a function, by means of a Turing machine or equivalent, is a strenuous exercise, and if all functions had to be strictly checked as to their computability in all mathematical activities, it would be impossible to have mathematics moving forward. This is where the **Church-Turing Thesis** enters the scene: "*Any effectively computable function is a recursive function - or Turing computable, etc.*" The way mathematicians use this thesis, and they use it quite often, can be exemplified by the following quotation from Johnstone [1987, page 40], in the context of recursive functions defined by Minsky's register machines:

"Henceforth we shall assume [Church's thesis] in the following informal sense: if we wish to show that a given function is recursive, we shall regard it as sufficient to give an informal description of an algorithm for computing it, without verifying explicitly that this description can be translated into a [Minsky's] register machine program."

In reality, no mathematician does mathematics in a totally mechanistic or automatic manner, in the way of a Turing machine; most steps performed rely on 'peer understanding', that is, part of what is being defined, as an algorithm for instance, relies on information that is not being written down explicitly. In this sense, *no one does mathematics*, or more appeasingly, mathematics done by humans is not entirely precise. "In short, mathematics restricts itself to what mathematicians consider *effective procedures*" [Kampis, 1991 p473].

This is exactly what Johnstone expresses in the previous quotation, to avoid going through the extensive explanation of a Turing or register machine program, defining a given function, we shall use an informal algorithm for computing it. The Church-Turing thesis provides the "alibi" for this lack of pure mathematical consistency, and it does so because no one has yet found a hole in this "more human" mathematics; that is, no one has presented an effective algorithm, accepted by the mathematical community, which does not have a strict mechanistic definition through Turing machines or other equivalent computability theories.

Example: The function string subtraction, for strings of a one symbol only vocabulary, $V=\{a\}$, can be effectively or informally defined as:

the corresponding turing machine definition takes at least two pages and can be seen in Loeckx[1972, pp 37-39].
$$\text{ssub}(a^x, a^y) = a^{x-y}, \forall x, y \in \mathbb{N} \mid x \geq y \geq 0$$

For mathematicians, the Church-Turing thesis is understood as: "*Mathematics problems can be solved only by doing mathematics*" [Hofstadter, 1979, p561]; in other words, until a hole is found in current computability theories by defining a more general theory of computation capable of defining more functions than those, we can do strict mathematics using human mathematics. Of course, if such a new theory is ever found, no doubt mathematicians will enlarge the thesis to the

Church-Turing-"whoever" thesis, expand the concept of computability, and continue to do human mathematics without the hassle of the mechanistic counterpart. In this sense, the Church-Turing thesis can be seen as a *tautological alibi*, that allows mathematicians not to strictly compute mathematics, unless when they so desire².

It is only when we step outside the (semantically) closed domain of mathematics that the Church-Turing Thesis and the notion of computation become really controversial. In general, any reproducible causal (natural) process is considered an *effective procedure*; in other words, if a process is 'reasonably well defined', not in a mathematical sense of course, but in the sense that we can manipulate or control the physical causes that lead to the reproducibility of our process, then it is an effective procedure. The notion of computation itself, is extended to the idea of a formal system which is functionally equivalent to an observed natural process; in other words, if we can find a formal system and equate its symbols and rules to some (chosen) set of observables and causal relations in the observed process, namely input-output, and in this way predict the functioning of such process, then we have *embodied computation*. The Church-Turing Thesis is thus used to correlate the computing abilities of formal devices with the causal dependencies of natural processes: if we can causally reproduce some natural process (extended effectiveness) then it is computable and it is performing an embodied computation. As a corollary, if we believe the universe is causal, then it is computable.

"According to Church-Turing, all physically realizable dynamics are equivalent to computation [...] [Conrad, 1985]"

"Church-Turing Thesis, Microscopic Version: The behavior of the components of a living being can be simulated on a computer. That is, the behavior of any component can be calculated by a [recursive function] to any degree of accuracy, given a sufficiently precise description of the component's internal state and local environment. [Hofstadter, 1979]"

"Living systems are nothing more than complex biochemical machines [...] [Crick, 1967]"

"Church-Turing Thesis, Reductionist version: All brain processes are derived from a computable substrate. [Hofstadter, 1979]"

"Church-Turing Thesis, AI Version: Mental processes of any sort can be simulated by a computer program [...] [Hofstadter, 1979]"

"Either Turing's thesis that effective procedures are computable is correct, or else consciousness may depend on hitherto unknown effective procedures that are not computable [...] [Johnson-Laird, 1983]"

Now, these extensions are indeed based on a leap of faith, and when using them one should be

² Several possibly effective algorithms have been proposed for which a computational description is not possible, see for instance [Kampis, 1991] for a good debate on this edition; however, ultimately for them to classify as holes in the thesis, they first have to be accepted as effective procedures which can always be denied, and thus render this problem undecidable.

conscious of its risks. The Church-Turing thesis was equated strictly within mathematics (even if its left hand side is not a strict mathematical concept), and extending it beyond those domains should be an act supported by the same kind of evidence. The fact that so far no one has presented an accepted effective algorithm that cannot be defined by an existent theory of computation, should not serve for evidence to the idea that since no one has proven that we cannot, in principle, find a formal model for any natural system, then all natural systems are computable. Unlike what happens in the mathematical domain, where every effective procedure found so far is (or can be) formally computable, in the natural domain it is very easy to find (extended) effective procedures, which are entirely causally reproducible, but for which there are not corresponding computational models, even in the extended sense.

The Church-Turing thesis cannot, honestly, be used in other domains other than mathematics, until someone shows evidence that *any* (extended) effective procedure found so far, *can be* described by a (extended) computational model, and not the other way around. In other words, the mathematical evidence for the thesis is not valid outside the mathematical domain closure; the thesis, in its extended sense, should not be used to claim anything, but simply to state a belief that natural causality can be described by formal descriptions. This belief does not find evidence in the generality of mathematical computability, and is as valid as any other belief, including its negation. Those that decide to work within the *Church-Turing belief system*, must not leave to 'the others' the task of disclaiming the thesis, they have to substantiate it themselves, if communicating in a scientific environment is the goal of their endeavors.

There is a circular argument in all the extensions of the Church-Turing Thesis beyond mathematics; the very first assumption that physically realizable dynamics represent effective procedures is just what the thesis aims at hypothesizing, in other words, the thesis is most often used precisely as the strongest basis for the assumption itself. This would not be a problem if the thesis was accepted by what it is (outside mathematics), an hypothetical belief; however, since all of its evidential weight is based on mathematical grounds, which have nothing³ to do with the physical world, claiming any such evidence for the extended thesis is confusing the simulation tool of some objects with all existing objects.

It is not the objective here to disqualify the Church-Turing Thesis belief system; I may even believe that formal mechanisms can, in principle, represent all of a causal physical universe, that is not the issue. I simply want to start this discussion by maintaining two basic points:

- The notion of computation is strictly a mathematical affair; good simulations of physical systems do not prove neither that the systems themselves are embodying some sort of computation nor that all physical systems are computable.
- Even if a certain level of a physical system is rendered a good enough

³ I am not forgetting the fact that mathematics must be at some primary level based on the physical world, or better, in some non-formal description of the physical world. Still, operations performed at higher levels of the mathematical edifice only very loosely, if ever, maintain the lower level semantic referents.

computational description (a good simulation), higher levels of this same system may not be appropriately characterized by the same description, as emergent properties may arise demanding descriptions based on their own terms.

The second point is not directly linked to the Church-Turing Thesis, however, and as seen in the previous quotations, there is a tendency to use the thesis as a motivation for reducing the functional levels of some system we wish to understand, to lower levels where basic components can be adequately simulated, and then render the whole system simulated. In the following, I will try to justify the belief that we should attempt to compare models at the functional levels we are interested in, even if a computational description is not there possible, as the knowledge of simulation of components may add nothing to the very same function we wish to understand.

2 - Computation and the External Observer.

The following definitions are needed for the subsequent discussion and as an expression of the position here maintained: A **Computation** is a dynamically incoherent (rate-independent) process, programmed in sequential steps which constitute the rule-based syntactic manipulation of small cardinality sets of discrete symbol vehicles. A **Formal System** is a set of symbols and a set of rules that manipulate such symbols in computational processes. A **Measurement** (and control) **process** is instead a dynamically coherent (rate-dependent), nonprogrammable activity, where only the end result is an observable or control constraint. If a formal system is used to model some real-world process, then a **modelling relation** is established between symbol vehicles in the computational processes and records of measured or control events to be executed in the process being modelled; these events are the contents that symbols *stand for*. The *symbol manipulation is syntactic* and hence independent from the semantic contents, this is precisely what allows the formal manipulation of symbols to be programmed. Notice that while contents are events constrained by dynamical time, their correspondent symbols are free from the grip of real time and can be manipulated by rate-independent computation.

One can of course, in principle, obtain a formal description for the measurement process, in other words, simulate the measurement; but to provide such simulation with values for its symbol vehicles, namely its boundary conditions, one has to perform measurements in the material world once again, and therefore, a completely formal description of reality is impossible as we encounter an infinite regress. This is known as *the measurement problem* (see [Pattee, 1986, 1989, 1992, 1993] [Von Neumann, 1955]).

"The basic measurement problem is that semantic grounding of symbols by measurements is a controlled action by an observer that cannot usefully (i.e. functionally) be described by laws. More precisely, if a measuring device -- which is certainly a physical system obeying laws -- is actually being described by these laws by combining the device with the original system being measured, then the initial conditions are no longer separated and additional new measuring devices are required to establish the initial conditions for this combined system. Therefore, syntactic description of measurement by laws destroys the semantic function of the measurement, and leads to a well-known infinite regress." [Pattee, 1993, page 115].

Underlying all of this is what one may call *the dogma of the perfect observer* [Medina-Martins and Rocha, 1992]; any formal system relies on an external observer/creator who ascribes meanings to its symbols by performing some form of measurement. In this respect a formal system can only participate in an open affair of the form *measurement* → *computation* → *control*, and can only be closed with the existence of an external observer 'feeding' its semantics. In contrast, a natural process does not have an external observer giving meaning to its symbols and dictating its next states; unless, of course, we bring back God into the picture. This is not obviously the idea of those working under a computationalist or functionalist approach; however, it is the disregard for the strict definitions of computation within mathematics, with all of what formal systems entail, that brings about the careless and undefined notions of embodied computation we are faced with today in fields such as AI and ALife:

"According to those AI theorists who take the human computational system as their model, there need be no difference between your computational procedures and the computational procedures of machine simulation, no difference beyond the particular physical substance that sustains those activities. In you, it is organic material; in the computer, it would be metals and semi-conductors. But this difference is no more relevant to the question of conscious intelligence than is a difference in blood type, or skin color, or metabolic chemistry, claims the (functionalist) AI theorist." [Churchland, Paul 1988, page 120]

The danger of using such loose notions lies in ignoring a crucial distinction between a computation and a natural process, precisely that of the external observer-God. In a physical device we decide to take as a computational tool, there is always an externally imposed relation between physical processes and the symbols in the computations we are performing, which is not a part of the device itself, but is in the mind some external observer. A physical computer only functions as such because its engineers decided to ascribe to certain physical processes, values such as zeroes and ones which can then be used as the basis of a hierarchy of more complicated symbols; but, this relationship is entirely external to the physical processes in the device. Because those engineers can control and reproduce those same physical processes, we obtain an extraordinary flexibility to manipulate them, and so, the symbols linked to those processes by an external relation can be very efficiently manipulated. Nevertheless, one must not forget that *the physical processes taking place do not depend at all on those symbols*, the control of the causes triggering such processes lies entirely on some external engineering. Thus, the computation is not embodied and lies in the minds of the engineers and users of such devices.

An example of a physical computer is the chinese abacus; its usage is based on ascribing certain patterns (of material beads) to numerical symbols, and devising rules to manipulate those bead patterns. Again, the computation is taking place in the mind of the user, and not on the device. Many other devices can be thought of; the degree of utility of such devices is proportional to the degree of control we possess on the material causes of the physical processes used. However, the computational processes they are aids for, remain in the minds of their creators and users.

An interesting case where this is apparent is the abacus-like device that the Inca scribes used in pre-colonial times. This device consisted of complicated patterns of knots in ropes; it is usually accepted that it was used somehow to maintain records of dates and as an aid in calculations;

however, many think that it was more than that, and it was capable of keeping entire narratives of events. Of course, its utilization was totally in the mind of the scribes, and the physical system we are faced with today offers no clue at all as to the device's function.

Formal symbol systems have syntax separated from semantics, and can only be meaningful within a closure formed by symbols/rules and the external semantics attributed by its users. If we loose the semantic referents, that is, open the closure and loose the semantic function, the syntactic part is meaningless. Let me once again go back to pre-colonial America for yet another example; consider the Mayan writing system, for years and years, linguists and historians have tried to understand it by looking solely at its symbols. It was believed that the language was purely, or at least largely, iconic, and so the symbols carried with them self-contained meanings; only when this idea was abandoned, and (phonetic) semantic referents were looked for in the spoken language of present day mayans, in other words, only when attempts at closing semantic loops close to the original were tried, did the mayan writing system begin to be understood. Unfortunately, and probably, the entire system semantics will never be understood as those who devised it, its engineers, are no longer around. For a report on this story see Michael D. Coe's [1989] "Breaking of the Maya Code".

In dynamic natural systems, in contrast, there is no external definition of the semantics of its objects; there is no program. The dynamics are entirely dependent on the (self) organization of its material constituents. The outputs of such processes are not the result of a sequential logical calculation which is rate-independent, but, of the interaction of different physical processes which are highly dependent on time. *The physical processes taking place are totally dependent on the (internal) meaning of its material constituents.* It can be said that meaning is an internal affair of the natural system: it is closed in its meaning.

We are then faced with an interesting problem; the systems we wish to model seem to follow a closed organization where self-reference is the key word. From the cell to the mind we find what may be described as *organizationally closed* processes, to use the usual autopoietic terminology, or, in other words, *closed action patterns of production* [Kampis, 1991], where we find that amongst the products of the system are the productive operators themselves [Pask, 1992] (e.g. DNA → mRNA → proteins → enzymes used in the production of DNA, mRNA and ribosomes).

Conversely, the computational tools we have to simulate these processes are open in the above discussed sense. Thus the tendency, when faced with complex systems (in Rosen's [1991] connotation of systems with closed causal loops leading to non-fractional properties), to descend to component levels of description, where the open simulation approach is easier and offers viable results. The problem is of course that these components offer no clue into the complex system's function, in other words, as we recognize our components to be more complex, we tend to descend to even lower component levels, loosing what we were trying to model in the first place.

Now, and as maintained before, perhaps a computational description is possible for living and cognitive systems at the functional levels we are interested in, that is, at the levels where we can

perceive living and cognitive capacities. However, the computational approach, with its semantically open attributes seems, at the very least, to render intractable problems when describing such systems. It seems to be a good idea to devise theories aiming at building conceptual models with largely closed and self-referential characteristics analogous to those of the natural systems we wish to model. These models, however, are not being computationally described at the functional levels we are interested in, and therefore building them amounts to devising analogues or metaphoric models, which though not mathematically explaining their function, may offer insights into the nature of the originals⁴. This objective is pursued in what follows.

3 - Pattee's Semantic Closure.

A very puzzling problem when studying biological systems is the question of symbolic description. Von Neumann presented the idea that it is more advantageous for an automaton to be able to assemble another automaton from an already existing description, rather than having to analyze the organization of the automaton to be constructed in order to obtain its description each time the process is to be performed [Von Neumann, 1966, page 83]. Generalizing, and intuitively, there seems to be an advantage in having organisms performing measurements on certain aspects of their organization and environment, their *boundary conditions*, and abstract them into some sort of code which can be manipulated in place of these same organizational aspects - a process of obtaining representations of boundary conditions.

In Von Neumann's self-reproducing automaton scheme the idea is precisely to have two different forms of manipulating the same basic automaton structure: uninterpreted and interpreted. There is a description of the self-reproducing automaton which is separately copied (uninterpreted or syntactically manipulated) and used for the construction of a new automaton (interpreted or semantically manipulated). This is sometimes referred to as a mixed active and passive mode of reproduction [Kampis, 1991], and it can be seen to refer, more generally, to Von Neumann's belief that natural organisms use matter in two distinct, complementary, modes:

"[Von Neumann's] most general conclusion was that natural organisms are mixed systems, involving both analog and digital processes. [...] In complicated organisms digital operations often alternate with analogue processes. For example, the genes are digital, while the enzymes they control function analogically." [Arthur Burks in Von Neumann, 1966, page 22]

"[...] processes which may go through the nervous system may, as I pointed out before, change their character from digital to analog, and back to digital, etc., repeatedly." [Von Neumann, 1958,

⁴ An interesting case of a conceptual model which is not really computable, is that of Von Neumann's [1966] self-reproducing scheme. The division between a machine and a description of a machine which must be handled separately on the reproduction process, offers great insights into the nature of biological reproduction. However, it is not a computable model, as there must always be a larger formal system used to define (simulate) the (Turing) machines involved in Von Neumann's automata, which is not reproduced with those machines. A formal system (defined by a Turing machine) cannot simulate itself; this fact is the basis of the famous halting problem. [See also Kampis, 1991, chapter 7]

By digital Von Neumann meant an all-or-none response pattern, and by analog a continuous physical response. Of course the strong dichotomy between digital and analog modes was from the start taken under a functional perspective, consider the following remarks in the context of nerve cells:

"[...] I have been talking as if a nerve cell were really a pure switching organ. It has been pointed out by many experts in neurology and adjacent fields that the nerve cell is not a pure switching organ but a very delicate continuous organ. In the lingo of computing machinery one would say it is an analog device that can do vastly more than transmit or not transmit a pulse. There is a possible answer to this, namely, that vacuum tubes, electrochemical relays, etc. are not switching devices either, since they have continuous properties. They are all characterized by this, however, that there is at least one way to run them where they have essentially an all-or-none response. What matters is how the component runs when the organism is functioning normally. Now nerve cells do not usually run as all-or-none organs. For instance, the method of translating a stimulus intensity into a frequency of response depends on fatigue and the time of recovery, which is a continuous or analog response. However, it is quite clear that the all-or-none character of a neuron is a very important part of the story." [Von Neumann, 1966, pages 68 and 69]

Though digital may have been too strong a term, especially in the context of neuronal activity of which we have a better idea today, it seems to me that the conception of natural organisms as mixed systems, where essentially continuous processes are used to hold an essentially discrete function, is very relevant. What is at stake here is the utilization of physical processes by other physical processes, not for the inherent physical characteristics of the first but for the ability of the second in manipulating those characteristics in a *meaningful* way. By meaningful manipulation I mean that the level of *user* physical processes manipulates the *used* physical processes in a manner relevant, or with significance, only to itself. A case of *matter using matter*.

It is precisely a two-level process like this that can be envisioned to hold and manipulate the representations of boundary conditions referred above: certain physical processes are used by other physical processes to represent specific boundary conditions relevant to the second. Loosely, we have the *user level* performing measurements on relevant aspects of its physical domain, and using processes in the *used level* to represent those measurements.

If we consider that a symbol is something *standing for* something else in some process, and which also presents discrete characteristics (in the sense that it can be utilized at any particular time, defined by some particular physical characteristics, without losing its specific *standing for* or semantic function), then those representations of boundary conditions can be regarded as symbols. It is important to emphasize that these symbols are not independent and self-evident entities, that is, they are only meaningful in the particular material structure that uses them. Further, these symbols are not manipulated by some external program, but by processes inherent to the closed structure where they can be considered meaningful symbol vehicles.

Howard Pattee's notion of *Semantic Closure* [Pattee, 1982, 1986, 1993], refers precisely to the

irreducible closure formed by measurement, control, and symbolic instruction actions, considered essential for the existence of evolvable biological or cognitive autonomy, i.e. Life. In order to secure the necessary functional and self-reference characteristics of living and cognitive systems, there must be an irreducible function-matter-symbol interdependence, escaping the computational external observer and dynamical incoherency.

This notion of closure is, in a sense, similar to that of Autopoiesis [Varela, 1979; Maturana, Varela, and Uribe, 1974] or Pask's Conversation Theory [Pask, 1976]; more generally, its most important feature is that of self-reference. The objective is the understanding of natural systems that observe autonomy, and for which the open computational modelling approach is very difficult, if not impossible. The important distinction between Semantic Closure and Autopoiesis, is that the first specifies the necessary closure as maintained in two levels, one material and the other symbolic; while the second specifies the minimum closure as being strictly a material affair. The principle of Semantic Closure considers the existence of symbolic characteristics of matter as a necessary part of any organizational closure responsible for biological and cognitive life. Even though the symbolic level of interactions can be reduced to a material level of organization, disregarding this symbolic organizational dimension misses a very important aspect of living organisms. The preferred example is that of genetic reproduction; DNA though with its own proper material characteristics, possesses a symbolic function extremely relevant for the understanding of the whole reproductive process, and which may be lost under a purely material analysis of molecular interactions.

When using the concept of symbol in the above sense, I do not mean to restrict myself to formal symbols, in which syntax (symbol manipulation) and semantics (symbol interpretation) are absolutely independent, and hence allowing programmable manipulation [as presented in sections 1 and 2]. The existence of natural symbols, whose syntax is not strictly separated from semantics since the symbolic level of interactions represents an emergent organizational level based on (and thus not separated from) a material organization, seems to be a relevant way of expressing processes taking place from the genetic to the natural language level, and perhaps even shed some light as to the invention of language itself, and the much more recent formal languages. [See Howard Pattee's contribution for this edition, for a more extensive presentation of his Semantic Closure principle, and also Charles Henry comments on the inseparability of syntax and semantics in natural language]. This same notion resonates in George Kampis' [1991] proposed Component-systems:

" Component-systems, which are inherently discrete, perform the same categorizing and chunking operations as we do in our discretized descriptions, and therefore they 'dcretize themselves'.

Now, if there are such systems, logic and discrete categories are no more just modes of description but modes of operation for them (cf. the 'linguistic modes' of Pattee). It is in the origin of such systems where we can recognize an answer to the question we asked when speaking about the origin of logic." [kampis, 1991, page 400]

The natural symbols existing in the context of a semantic closure, are not allowed programmable manipulation (computation) in account of their not strictly separated syntax and semantics. The

internal symbol manipulation by semantically closed systems is sometimes referred to as *internal programming* [Kampis, 1993]. Though intuitively this may be a useful concept, it is also a misleading one, since internal programming is defined as non-algorithmic and non-computational; this, I fear, encourages the popular metaphoric use of the notion of computation, strengthening the qualification of everything as computational processes by the predominant computationalist attitude in current scientific activity. Much like putting out fire with gasoline. In any case, it is a notion imbedded in the spirit of what was discussed above:

"[...] Externally programmed systems have externally supplied semantics which in the system appears as part of the syntax (hence not meaning-laden at all), a known paradox of logic and formal systems. Internal programming can, on the other hand, establish own semantic relations. So, whereas external programming is rigid and a priori, internal programming can be soft and dynamic.

Perhaps this is a point where Searle and computers can be brought nearer to each other. Searle's criticism applies for external but not for internal programming, the latter being non-algorithmic, and hence capable of constructing its own terms. The internal programming actions when performed generate a meaning that is interpretable within, but not outside the system. This is exactly the opposite of the usual computational situation where a hard-wired meaning allocation excludes the system from manipulation and creation of its meaning, leaving it to an outside intelligence, the programmer." [Kampis, 1993]

A final word in this section regarding these natural symbols. Von Foerster [1977, 1981], introduced the concept of *eigen-value* as values (functions, operators, algorithms) that could satisfy an indefinite recursive equation describing the organization of sensori-motor interactions, or more generally self-referential organizations such as cognitive systems. He concluded first, from the possible existence of such values (with no true mathematical description), that they must be discrete even if their primary argument or observable is continuous, since in a recursive loop where they are generated only stable observables will maintain eigen-value representation:

"In Other words, Eigenvalues represent *equilibria*, and depending upon the chosen domain of the primary argument [measured observable], these equilibria may be equilibria values ("Fixed Points"), functional equilibria, operational equilibria, structural equilibria, etc." [Von Foerster, 1977, 1981, page 278].

He further concluded that because of their self-defining nature, eigen-values imply a "topological closure" or circularity. The idea is thus that eigen-values are discrete, stand for stable observables, and can only exist within an operational closure, in other words, eigen-values are symbols standing for measured boundary conditions, available and meaningful only within a functional-matter-symbol interdependence that can be referred to as semantic closure.

I like to consider an image "grabbed" from Von Foerster [1981] (where it is presented under a somewhat different context), as a visual metaphor for semantic closure. The two snake-dragons eating each other's tail can be seen as the two sides of matter forming a semantic closure: discrete/continuous, symbolic/physical, syntactic/semantic - closed and inseparable. Matter using

Distinctions between the two paradigms must be understood as functional distinctions at the level of description we are interested in, and not in terms of formal distinctions based on computational simulations, that may be possible at some particular level, but which nonetheless are syntactic and devoid of any functional meaning at the relevant level. The bottom line is that knowledge representations in the formal models are accessible and alterable from the standpoint of the external observer of the computation; while in the connectionist models, the distributed and superposed memory scheme [Van Gelder, 1992], makes those representations nowhere to be found and therefore not explicitly alterable: there is no external computation. The two paradigms are indeed *horses of a different color*. Further, if formal descriptions are the only accepted models for natural systems, then connectionist devices cannot be considered models of memory or cognition, but simply models of the neuron. An alternative might be to label such devices analogues or *metaphoric models* of memory. Since it is beyond the scope of this essay to argue for the scientific legitimacy of non-formal models, I will simply state my intent of construing such models when appropriate based on my belief that they can add to scientific knowledge in a positive way. Notice, for instance, that there is no formal model for evolution, yet it is indeed a genuine scientific concept.

But are Artificial Neural Networks (ANN's) a panacea for our cognitive modelling problems, will everything come up roses as we accept current connectionism as a more viable approach to cognition? Clearly, these models have already proved remarkably well in an ever growing number of applications, enough to put Neural Networks way up on the *Buzz Word Hall of Fame*, however, as models of the brain, their current characteristics seem to be a mere scratch on the surface.

These models are functionally nonprogrammable, however they do not offer yet dynamically coherent processes, and in that respect seem to fall in a strange category of *nonprogrammable computational* or *dynamically incoherent measurement devices*. Both of these labels are paradoxical unless, of course, we decide to enlarge either the notion of computation or the notion of measurement. Non-programmable computation seems to be ambiguous to such a degree that the term becomes meaningless; enlarging the concept of measurement seems to be the best bet. However, for the objectives of the current work, we need not worry about this problem, as it is its objective to lay the grounds for building connectionist devices which *are* dynamically coherent, and hence measurement devices in the strict sense.

We have already seen that ANN's are nonprogrammable models, which is a result of a certain level of closure of the organization of their internal constituents: the artificial (computational) neurons. The knowledge representations connectionist systems keep are a closed, internally only meaningful affair. But, are they semantically closed, or completely self-referential artifacts?

The neuronal level of interactions, (self) organizes only with the inclusion of inhibitory and excitatory signals received from the outside: from an external observer judging outputs from inputs in a learning process. Hence, the system as a whole is not closed; to obtain the necessary closure, these inhibitory and excitatory signals must be included in its organization, without explicit external action, as self-states, or, using Von Foerster's [1969] term, as *eigen-states*.

Perhaps an interesting idea to pursue is to consider them as symbols in a broad sense discussed earlier, in which case they would allow the formation of the symbolic level of this proposed semantic closure. These symbols would of course be inseparable from their structural implementation and would only be meaningful inside the closure.

Notice that considering such signals as symbols, amounts to the *broader-than-formal* definition of symbol discussed before: symbols as representations of stable localized boundary conditions measured in the material level of the organization, and utilized in an emergent symbolic level. Conversely, in strictly material approaches to biological and cognitive systems, such as autopoiesis, the interactions caused by the would-be symbol vehicles (or representations of boundary conditions) are considered simply a part of the necessary material organizational closure. Following my remarks on previous sections, I pursue the idea that failing to recognize the emergence of symbolic properties in certain material structures, may take the research on living organizations down to material levels of interaction which may offer no insight into the very function we wish to understand.

"[...] I just want to point out that the componentry used in the memory may be entirely different from the one that underlies the basic active organs. [Von Neumann, 1958, page 67]"

5 - Von Foerster's Cognitive Tiles.

"[...] We know the basic active organs of the nervous systems (the nerve cells). There is reason to believe that a very large-capacity memory is associated with this system. We do most emphatically *not* know what type of physical entities are the basic components for the memory in question." [Von Neumann, 1958, page 68]

It was precisely in an attempt to develop a model of memory that Von Foerster [1969] proposed his *Tessellations of Cognitive Tiles*. The idea was to advance a conceptual minimum element capable of observing the desired cognitive characteristics of memory; this model, though using connectionist devices, was not identified with a neuron or any particular neuronal arrangement to be found in the brain, but simply proposed as a conceptual model of the desired memory function. By including a self-referential level of interactions, in which an internal meaning of (measured) states of the memory-empowered organization is generated (the eigen-states), a closed system, with internal meanings, is obtained, "so that we may ultimately dismiss the demon [the external observer] and put his brain right there where ours is" [Von Foerster, 1969]. Let us now look into a *Cognitive Tile* in more detail.

A Cognitive Tile, as seen in Figure 2, is a through-put system. Sensory information, \mathbf{X} , is compared, in a feed-forward fashion, and altered in respect to the self ($\mathbf{R}_s(\mathbf{X})$). All information is checked as to its usefulness by comparison with the tile's eigen-states, in other words, with regard to its relevance to the tile's function, expressed in the symbolic attributes of the closure. Consider for instance that the objective of one of these tiles would be to maintain the level of a specific chemical within some interval; all inputs arriving to the tile are then altered according to their significance regarding the tile's current eigen-states, namely representations of the (measured) desired chemical level and measurements of the current structure, which under such a perspective start to hold symbolic attributes. The symbol vehicles are elements of some material structure available for utilization by the symbolic level of the closure.

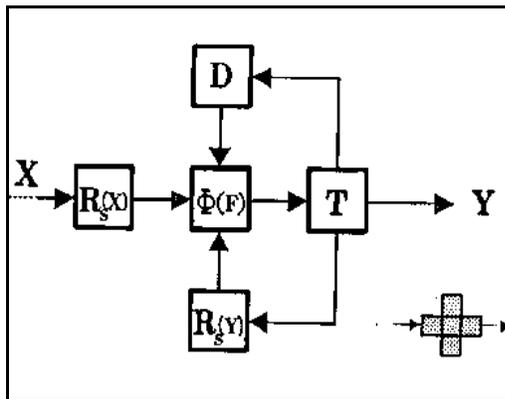


Figure 2: Circuitry and basic components of a cognitive tile.

\mathbf{X} - Sensory Information;
 $\mathbf{R}_s(\mathbf{X})$ - Self-Referential Block. Incoming Information is checked as to its usefulness.
 $\Phi(\mathbf{F})$ - Central Block. Distributed Memory Complex. Artificial Neural Network.
 \mathbf{T} - Transducer. Translates Output into signals \mathbf{Y} that are understood by components it communicates with, including the external world.
 $\mathbf{R}_s(\mathbf{Y})$ - Self-Referential Block. Output Information is checked as to its usefulness. This dictates the new self-states which will affect the next organization of the

The compared (and altered according to their significance) inputs are then fed to the central block defined as a distributed memory complex (realized by some ANN), implementing, at a specific time t , a particular classification "function" \mathbf{F}_t which is not being formally described (or computed). It is best described as a set of *material implications*, linked to the very fabric of their implementation, allowing us nonetheless to write $\mathbf{Y}_t = \mathbf{F}_t(\mathbf{X}_t)$; in other words, the output \mathbf{Y}_t is uniquely obtained from the input \mathbf{X}_t as it "finds" the material function \mathbf{F}_t . Then, similarly to any connectionist model, the material function \mathbf{F}_t will be altered at each cycle when the report of the effect of the output \mathbf{Y} is available. This can be conceptualized as we consider \mathbf{F}_t to be a range of possible material functions in the domain of a *material functional* $\Phi(\mathbf{F}_t)$. The notation is merely a mathematical formalism metaphor, as the material functions are not formally defined.

This conceptualization serves only the purpose to indicate that the system acts according to the outcome of previous actions, the present output is dependent on the history of the previous inputs, and ultimately, on the initial material function \mathbf{F}_0 . In this sense it is a state-determined system, but not in the traditional formal sense, since the material functions \mathbf{F}_t are not being computed at any stage; there is no program, only when inputs "find" the present material function will the output be available. *Outputs are not computed at this level of description.* This is the

way any ANN functions, and it should not be confused with the fact that the material functions may in principle admit a formal simulation, or that at a different level of description ANN's are formally defined and therefore computable. At such a neuronal level they lose their memory function and become meaningless: a weight or component analysis there performed, will not be able to retrieve any memory contents⁶. Memory does not exist in the computational level of weights and adding machines, it is everywhere, unretrievable, in the present state of the measurement device working at the functional level we are interested in.

It may help to recall here Gordon Pask's [1958, 1959, 1960] 1950's electrochemical device which constructed a form of self-organizing network (under a reward scheme), adaptively implementing some functional task like sensitivity to sound. Such a network rather than a computational lower level of description had an electrochemical functional description, which nonetheless was also unable to describe the upper level's function properly. [See Cariani, 1989, 1993.]

Going back to the characterization of the Cognitive Tile's organization; the central block will then output a particular action. The results of this action may alter the eigen-states, and therefore the (several) next material functions, through the lower (self-related) feedback loop $\mathbf{R}_s(\mathbf{Y})$. \mathbf{T} is a transducer, which translates the output of the central block \mathbf{Y} into signals that are understood by the several components it communicates with, including the exterior world. Another way of looking at this block, is to consider it a measurement block evaluating the boundary conditions of the material level, into eigen-states operating in the symbolic level of the closure. This last interpretation of the transducer block, pursued here, demands the introduction of a similar block in the feedforward link to be proposed ahead. The symbolic level is thus enabling the communication between two distinct systems: the subsymbolic memory complex existing in the central block, and the physical organization present where measurements are performed.

The upper feedback loop simply incorporates the delay \mathbf{D} of the system, that is, the time-scale of the Tile. This time-scale is the clock that dictates when an evaluation of \mathbf{F} regarding the eigen-states should be performed, in order to assess its earlier actions. It is also a part of the symbolic level as it functions on information supplied by the transducer-measurement device \mathbf{T} .

This element includes the faculties necessary for maintaining semantic closure. There is an irreducible function-matter-symbol interdependence: (memory) function is established on the actions of the tiles which are a result of the history of all previous input-action pairs and the corresponding self-referential reactions, all of which cannot be abstracted from its material implementation (even if it turns out to be computational); the transducer \mathbf{T} implements the connection between the material and symbolic levels, as it measures the material boundary conditions at each cycle, converting them into symbols that are operated by the self-referential block; the self-referential block will, from the symbols received from \mathbf{T} , trigger the appropriate actions to alter the central block's neuronal aggregate and thus affect its future functional response. Hence, there is an irreducible closure formed by measurement, control, and symbolic

⁶ Except in some special simpler network cases.

instruction.

We can think of these tiles as a suggested conceptualization of the necessary connections between symbol and matter, in order to obtain an autonomous classification function or semantic closure. This conceptualization may still be far from complete, as it is an on-going research interest, therefore only by actually building and testing the Cognitive Tiles will we be able to assess both their value as classifiers and their degree of autonomy. In the next sections I will try to lay the very initial foundations for building models using these concepts.

6 - Using and Expanding the Tiles: Dynamic Neurons.

Cognitive tiles can be incorporated into large tessellations; initially Von Foerster stated that information exchange between tiles can take place on all interfaces, as long as we observe the flow diagram of figure 2. It is not really clear which communication channels are allowed in these original tessellations; moreover, to make sure that information is only taken in reference to the tile's own function, all information must get into the tile exclusively through the first self-referential block, so that it is checked and altered as to its usefulness. Since the self-referential blocks operate on the symbolic level, another transducer-measurement device T_{in} must be included in the beginning of the feedforward link. This new block will measure the boundary conditions of the organization in the beginning of each cycle; the original transducer, now dubbed T_{out} , will measure boundary conditions at the end of each cycle. Figure 3 presents this slightly different version of a cognitive tile.

Notice also the inclusion of two extra information channels: the upper one is a feedforward control channel which enables inputs (and the measurement of boundary conditions) to affect the time-scale D of the element; the bottom one is actually a two-way channel which indicates the fact that both self-referential blocks must affect each other; or better, are part of the same self-referential function. However, it is important to observe that the feedback self-referential block, $R_S(Y)$, recognizes current eigen-states and influences the central block accordingly, while the feedforward self-referential block, $R_S(X)$, recognizes these same eigen-states and influences not the central block, but information supplied to the central block.

The implementation of each of these blocks presents us with the first batch of interesting questions, problems, and possibilities for modelling these autonomous classifiers. Allow me though to regress just a bit to the objective of a cognitive tile itself; there are two ways to

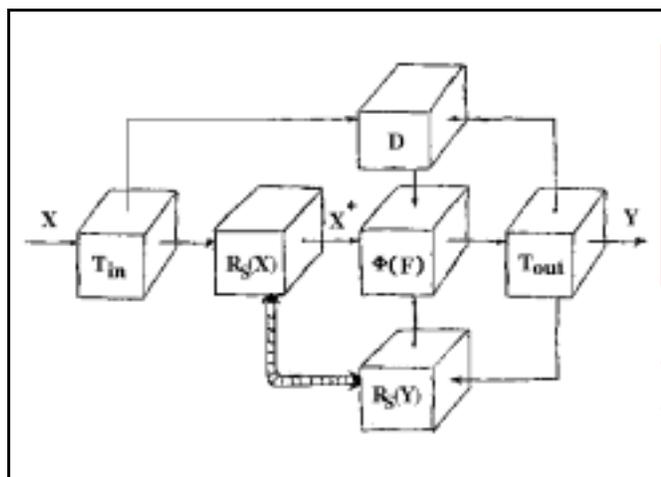


Figure 3: Expanded Cognitive Tile

start approaching this issue: either we accept the existence of symbols as a primitive of the system, in which case their specific materiality is beside the point, or, instead decide to study the very self-organizing processes that may lead to the arising of internal symbols. The first approach works with cognitive tiles more as a conceptualization of the symbolic part of a semantic closure, while the second tries to study precisely the emergence of this semantic closure from material (or simulated material) constituents. Undoubtedly the second problem is the most interesting⁷, it is not however what is pursued at this time, and rather something to hopefully address at a later stage. Here I simply want to lay the grounds for a possible utilization of semantically closed tiles as building blocks for more dynamic models of memory. Figure 4 shows a more detailed look at these semantically closed primitives we wish to assemble.

The black and white circles in the figure represent measurement and control devices; if the flow is from black to white it means a measurement (or material to symbolic transfer) is performed, the opposite flow means a control action (symbolic to material transfer) is performed. Lowercase characters and curved lines represent material flows, while uppercase characters and straight lines represent symbolic flows. In the light of the previous remarks, it is accepted as a primitive that the measurement devices transform aspects of the material organization and environment into symbols that the symbolic blocks can handle, and which represent those boundary conditions. It is further accepted, at least as a first tactic, that the self-referential blocks are biased towards some desired stability regarding the boundary conditions.

The tiles to be built can be based on a series of alternatives. The central block can be implemented by a ANN with various degrees of complexity, from a simple network of perceptrons, to backpropagation, Hopfield, recurrent networks, or even temporal pattern

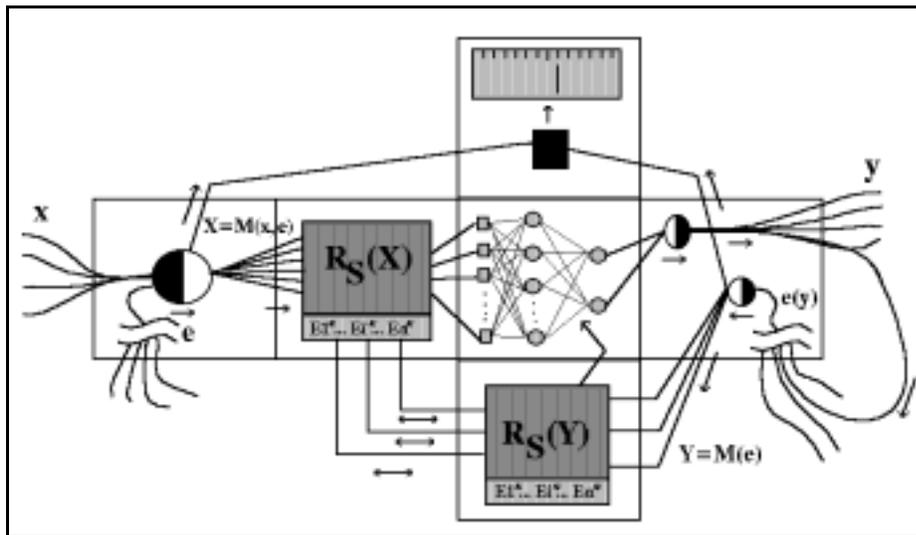


Figure 4: A more detailed Tile

classifiers [see Peter Cariani's essay on this subject in the present edition]. The self-referential and time-scale blocks can be implemented by hardwired programming, in which case, we are partially opening the closure and thus programming the desired self-referential characteristics of the tiles, or through the

⁷ Very related to the problem of the origin of life, see Howard Pattee's presentation on this same edition.

use of ANN's⁸. The tiles I am currently studying are composed of simple perceptron networks in each of these blocks, and with a small number of symbolic inputs. The first objective is to obtain control of the internal time-scale, in other words, obtain a causal relationship, between patterns of inputs and variations of time-scales.

It seems to be a good idea to use these elements not only in higher level modelling approaches to memory and cognition, as traditional ANN's, but also at the presumable lower level of Artificial Life models. For example, in a cell we can find different processes with different time-scales, if these processes can be organized into semantically closed groups, then they can be represented by Cognitive Tiles, and the functioning of a cell by a tessellation. A simple example of an artificial cell can be seen on Figure 5. Notice that we have processes that affect the time-scales and self-states of others. With an arrangement such as this, we may be able to obtain true temporal pattern recognition and not simply sequence recognition

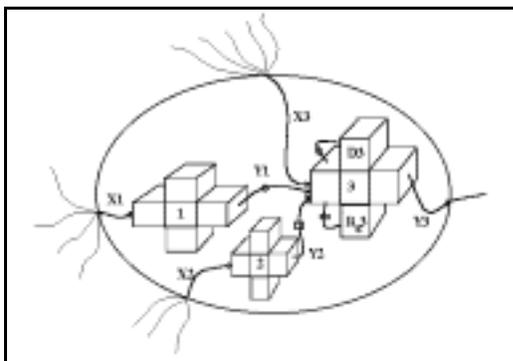


Figure 5: Artificial Cell with three tiles.

as is the case of recurrent ANN's. We can thus start to consider a tessellation of cognitive tiles as a proper measurement device, which, even with a computational lower layer, becomes dynamically coherent

Hence, why not model a single neuron through such a tessellation? Current research in Neurophysiology [Churchland and Sejnowski, 1992] clearly indicates the complexity of these cells which should not be modelled by memoryless response functions. They can be considered as parallel distributed processors implemented by a network of interacting electrochemical systems; these processors will be responsible for the recognition of temporal, rather than spatial, patterns on the array of inputs. The recognition of appropriate temporal patterns will then dictate the neuron's response. Clearly, time plays the most important role on the functioning of these dynamic neuron's, unlike the time-less switches of traditional ANN's.

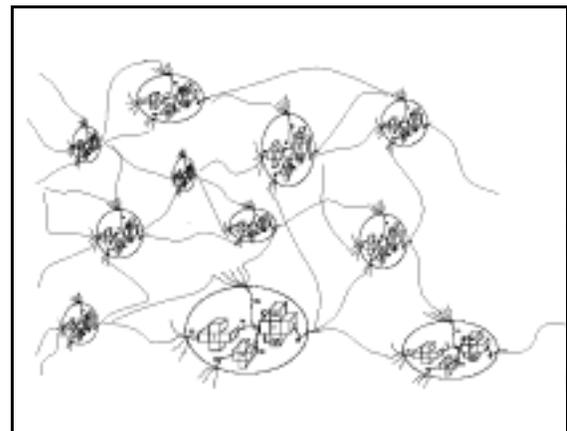


Figure 6: Network of three-tile neurons

⁸ It is through the use of ANN's in the self-referential blocks that we may attempt to study the problem left behind concerning the emergence of symbols themselves; by abandoning the constraint of bias to some desired boundary conditions (desired eigen-states), the only alternative is to use unsupervised networks that may learn to cluster incoming information into eigen-state representations in a way as to direct the central block's network to output appropriate actions leading to the survival of the tile in a competitive environment. It is a much more complex problem demanding the introduction of selective environmental constraints, though very relevant for the study of Artificial Life, and can only be regarded at this time as a future research interest.

We can think of networks where each neuron is represented by tessellations of cognitive tiles (figure 6).

As we understand better the complex networks of tessellations presented above, perhaps they can be organized as blocks of tiles dynamically closing higher level semantic loops based on other semantic loops. In the end, however, and as any such idea, all of this must be judged under a modelling perspective. We need to develop real models with these complex networks (of tessellations) in order to assess their value. It is such a modelling endeavor that I am currently pursuing, the on-going results of which will be left for a future presentation.

7 - Summary.

I have presented here a conceptual memory-empowered element observing semantic closure based on a slight extension of Von Foerster's Cognitive Tiles. These elements offer a new approach to the modelling of biological and cognitive systems, as they propose a semantically closed, non-computational building block to construct such models. Their first application was proposed in the development of more dynamic Artificial Neural Networks, where neurons are complex, rate-dependent, memory-empowered processors.

Acknowledgments.

Figures one and two were scanned directly from Von Foerster's [1969, 1981] original presentation of cognitive tiles and eigen-behaviors with his kind permission.

References.

- Cariani, Peter [1989], "*On the Design of Devices with Emergent Semantic Functions*". Phd Dissertation. Systems Science Department, Watson School, SUNY-Binghamton.
- Cariani, Peter [1993], "To Evolve an Ear: Epistemological Implications of Gordon Pask's Electrochemical Devices". *Systems Research* Vol.10 No. 3, pp. 19-33.
- Church, A. [1936], "An unsolvable Problem for Elementary Number Theory". *American Journal of Mathematics*, Vol 58, pp 345-363.
- Churchland, Patricia S. and T.J. Sejnowski [1992], *The Computational Brain*, MIT Press, Cambridge, Massachusetts.
- Churchland, Paul M. [1988], *Matter and Consciousness: A Contemporary Introduction to the Philosophy of Mind (Revised Edition)*. MIT Press, Cambridge, Massachusetts.
- Coe, M.D. [1989], *Breaking of the Maya Code*. Thames and Hudson, London, U.K..
- Conrad, M. [1985], "On Design Principles for a Molecular Computer". *Comm. ACM* Vol 28, pp 464-480.
- Crick, F. [1967], *Of Molecules and Men*. University of Washington Press. Seattle. USA

- Von Foerster, Heinz [1969], "What is Memory that it may have Hindsight and Foresight as well?". In *The Future of the Brain Sciences*. Ed Samuel Bogoch. Plenum Press, New York, pp 19-65 and 89-95.
- Von Foerster, Heinz [1977], "Objects: Tokens for (Eigen-) Behaviors". In *Hommage a Jean Piaget: Epistemologie genetique et equilibration*, Eds. B. Inhelder, R. Garcia, J. Voneche, Delachaux et Niestel, Neuchatel. Reprinted in Von Foerster [1981].
- Von Foerster, Heinz [1981], *Observing Systems*, Intersystems Publications, California.
- Hofstadter, D. [1979], *Gödel, Escher, Bach*. Basic Books, New York, USA.
- Johnsone-Laird, P.N. [1983], *Mental Models*. Cambridge University Press, U.K.
- Jonhstone, P.T. [1987], *Notes on Logic and Set Theory*. Cambridge University Press, U.K..
- Kampis, George [1991], *Self-Modifying Systems in Biology and Cognitive Science*, Pergamon Press.
- Kampis, George [1993], "Organization, Not Behavior: An Essay about Natural and Artificial Creatures". Paper delivered at the Workshop *Artificial Life: A Bridge towards a New Artificial Intelligence*, organized by the Department of Logic and Philosophy of Science, University of the Basque Country, San Sebastian, Spain, December 1993.
- Loeckx, J. [1972], *Computability and Decidability: An Introduction for Students of Computer Science*. Springer-Verlag, Berlin, Germany.
- Markov, A.A. [1951], "Theory of Algorithms". American Math. Society Translations. Vol 15 [1960], pp 1-14. (Translation of the Russian Original).
- Medina-Martins, Pedro R. and Luis Rocha [1992], "The In and the Out: An Evolutionary Approach". In *Cybernetics and Systems Research '92*. Ed. R. Trappl, World Scientific Press, Hong Kong, pp 681-689.
- Minsky, M. [1967], *Computation: Finite and Infinite Machines*. Prentice-Hall, Englewood Cliffs.
- Von Neumann, J. [1955], *Mathematical Foundations of Quantum Mechanics*. Princeton University Press. Princeton, New Jersey, USA.
- Von Neumann, J. [1958], *The Computer and the Brain*. Yale University Press.
- Von Neumann, J. [1966], *The Theory of Self-Reproducing Automata*, Edited and completed by Arthur Burks, Univ. of Illinois Press, Urbana.
- Pask, Gordon [1968], *An Approach to Cybernetics*, Hutchison Publishers, London.
- Pask, Gordon [1976], *Conversation Theory: Applications in Education and Epistemology*, Elsevier Scientific Publishing Company, Amsterdam.
- Pask, Gordon [1992], "Different kinds of Cybernetics." In *New Perspectives on Cybernetics*, Ed. G. Van de Vijver, Kluwer Academic Publishers, Netherlands, pp 11-31.
- Pattee, Howard H. [1982], "Cell Psychology: An Evolutionary Approach to the Symbol-Matter Problem".

- Cognition and Brain Theory, 5(4), pp 325-341.
- Pattee, Howard H. [1986], "Universal Principles of Measurement and Language Functions in Evolving Systems". Reprinted In *Facets of Systems Science*, George Klir [1991], Plenum Press, New York. pp 579-592.
- Pattee, Howard H. [1989], "The Measurement Problem in Artificial World Models". In *BioSystems*, Vol. 23, 1989, Elsevier Scientific Publishers, Ireland, pp 281-290.
- Pattee, Howard H. [1992], "The Measurement Problem in Physics, Computation, and Brain Theories". In: *Nature, Cognition, and Systems*, Vol 2, pp 179-192, Ed. M. Carvallo, Kluwer Academic Publishers, The Netherlands.
- Pattee, Howard H. [1993], "The Limitations of Formal Models of Measurement, Control, and Cognition". *Applied Mathematics and Computation* 56, pp 111-130, Elsevier Science Publishing.
- Rosen, Robert [1991], *Life Itself: A Comprehensive Inquiry into the Nature, Origin, and Fabrication of Life*. Columbia University Press, New York.
- Post, E. [1965], "Absolutely Unsolvable Problems and Relatively Undecidable Propositions". In M. Davis [1965] (Editor). *The Undecidable*. Raven Press, New York. pp 340-433.
- Turing, A.M. [1936], "On Computable Numbers, with an Application to the Entscheidungsproblem". *Proc. London. Math. Soc.* Vol 42, pp 230-265.
- Turing, A.M. [1937], "Computability and lambda-definability". *Journal of Symbolic Logic* Vol 2, pp 153-163.
- Van Gelder, T. [1992], "What is the 'D' in 'PDP'? A Survey of the Concept of Distribution". In *Philosophy and Connectionist Theory*. Eds. W. Ramsey, S. Stich, and D. Rumelhart, Lawrence Erlbaum Associates, New Jersey
- Varela, F. [1979], *Principles of Biological Autonomy*. New York: Elsevier North Holland.
- Varela, F., H. Maturana, and R. Uribe [1974]: "Autopoiesis: The Organization of Living Systems, its Characterization and a Model". *Bio. Syst.* Vol 5, 187.

AI	1, 3
ALife	1, 3
Artificial Life	6
Computability	1, 2, 1
Computation	1-4, 8, 1
Connectionism	4, 1
Measurement	2-4, 6-8, 1
Semantic Closure	1, 4-6, 8, 1
Systems	1